

(4)

Segment-based Acoustic Models for Continuous Speech Recognition

Progress Report: 1 July 94 – 31 December 94

submitted to
Office of Naval Research
and
Advanced Research Projects Administration

by
Boston University
Boston, Massachusetts 02215

Principal Investigators

Dr. Mari Ostendorf
Associate Professor of ECS Engineering, Boston University
Telephone: (617) 353-5430

Dr. J. Robin Rohlicek
Division Scientist, BBN Inc.
Telephone: (617) 873-3894

Administrative Contact

Maureen Rogers, Awards Manager
Office of Sponsored Programs
Telephone: (617) 353-4365



Accession For	
NTIS CRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification _____	
By _____	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

DTIC QUALITY INSPECTED 3

DISTRIBUTION STATEMENT A
Approved for public release; Distribution Unlimited

19950124 002

Executive Summary

This research aims to develop new and more accurate stochastic models for speaker-independent continuous speech recognition by extending previous work in segment-based modeling, by introducing a new hierarchical approach to representing intra-utterance statistical dependencies, and by developing language models that capture topic dependencies. These techniques, which have high computational costs because of the large search space associated with higher order models, are made feasible through a multi-pass search strategy that involves rescoring a constrained space given by an HMM decoding. We expect these different modeling techniques to result in improved recognition performance over that achieved by current systems, which handle only frame-based observations and assume that these observations are independent given an underlying state sequence.

With these overall project goals, the primary research efforts and results over the last two quarters have included:

- implementation of several software system improvements to enable research in more general distribution clustering and score combination weight estimation;
- development of a constrained EM algorithm for training the mixture language model which led to a small improvement in performance over Viterbi-style training;
- development of a mixture version of cache language modeling, together with a new content-word cache model, obtaining a small error reduction for short (3-sentence articles);
- implementation of the EM training algorithm for dependence tree design and experimental exploration of parameter smoothing;
- development of new approaches to channel compensation, based on a Bayesian approach of estimating a prior for the channel; and
- implementation and evaluation of three lattice search algorithms, providing an understanding of conditions under which the different algorithms are most appropriate.

We also participated in the November 1994 ARPA benchmark recognition tests, and achieved 11.6% word error rate using an old version of our system (since the new developments in our system were not integrated in time for the benchmark). This performance level is comparable to the mid-performance systems in the evaluation. We hope to be able to re-evaluate with more recent improvements early in the next quarter.

Contents

1	Productivity Measures	4
2	Summary of Technical Progress	5
3	Publications and Presentations	12
4	Transitions and DoD Interactions	13
5	Software and Hardware Prototypes	14

Principal Investigator Name: Mari Ostendorf

PI Institution: Boston University

PI Phone Number: 617-353-5430

PI E-mail Address: mo@raven.bu.edu

Grant or Contract Title: Segment-Based Acoustic Models for Continuous Speech Recognition

Grant or Contract Number: ONR-N00014-92-J-1778

Reporting Period: 1 July 1994 – 31 December 1994

1 Productivity Measures

- Refereed papers submitted but not yet published: 1
- Refereed papers published: 1
- Unrefereed reports and articles: 3 graduate student theses
- Books or parts thereof submitted but not yet published: 0
- Books or parts thereof published: 0
- Patents filed but not yet granted: 0
- Patents granted (include software copyrights): 0
- Invited presentations: 1
- Contributed presentations: 0
- Honors received: none
- Prizes or awards received: none
- Promotions obtained: none
- Graduate students supported $\geq 25\%$ of full time: 5
- Post-docs supported $\geq 25\%$ of full time: 0
- Minorities supported: 2 women

Principal Investigator Name: Mari Ostendorf

PI Institution: Boston University

PI Phone Number: 617-353-5430

PI E-mail Address: mo@raven.bu.edu

Grant or Contract Title: Segment-Based Acoustic Models for Continuous Speech Recognition

Grant or Contract Number: ONR-N00014-92-J-1778

Reporting Period: 1 July 1994 – 31 December 1994

2 Summary of Technical Progress

Introduction and Background

In this work, we are interested in the problem of large vocabulary, speaker-independent continuous speech recognition, and primarily in the acoustic modeling component of this problem. In developing acoustic models for speech recognition, we have conflicting goals. On one hand, the models should be robust to inter- and intra-speaker variability, to the use of a different vocabulary in recognition than in training, and to the effects of moderately noisy environments. In order to accomplish this, we need to model gross features and global trends. On the other hand, the models must be sensitive and detailed enough to detect fine acoustic differences between similar words in a large vocabulary task. To answer these opposing demands requires improvements in acoustic modeling at several levels: the frame level (e.g. signal processing), the phoneme level (e.g. modeling feature dynamics), and the utterance level (e.g. defining a structural context for representing the intra-utterance dependence across phonemes). This project addresses the problem of acoustic modeling, specifically focusing on modeling at the segment level and above. The research strategy includes three main thrusts. First, phone-level acoustic modeling is based on the stochastic segment model (SSM) [1, 2], and in this area our main efforts involve developing new techniques for robust context modeling, mechanisms for effectively incorporating segmental features, and models of within-segment dependence of frame-based features. Second, high-level models are being explored in order to capture speaker-dependent and session-dependent effects within the context of a speaker-independent model. In particular, we are investigating hierarchical structures for representing the intra-utterance dependency of phonetic models, and more recently language models for representing topic dependency and language dynamics, recognizing that higher-order models of correlation can extend to the language domain as well as the acoustic domain. Lastly, speech recognition is implemented under a multi-pass search framework, which in most of our work has been based on the N-best rescoring paradigm [3] where we use the BBN Byblos system is used to constrain the SSM search space by providing the top N sentence hypotheses. This paradigm facilitates research on high-order models through reducing development costs, and provides a modular framework for technology transfer that has enabled us to advance state-of-the-art recognition performance through collaboration with BBN.

Summary of Technical Results

In brief, the accomplishments of previous work on this project have included: improvements to the N-Best rescoring weight estimation algorithm; investigation of different mechanisms for improving the acoustic model, including distribution clustering [4], mixture modeling at different time scales [5, 6], theoretically consistent models based on context-dependent posterior distributions [7], automatic distribution mapping estimation, and hierarchical models of intra-utterance phoneme dependence [8]; development of a new approach to adaptation of continuous density parameters; and implementation of baseline n-gram and sentence-level mixture language models [9]. In addition, we have regularly participated in the ARPA speech recognition benchmark tests.

The research efforts during this period, supported in part by AASERT awards, have emphasized search and modeling techniques for long-distance knowledge sources, software development, and participation in the November 1994 ARPA benchmark test. These efforts and the primary research results are summarized below. It has been a busy period, evidenced by the fact that one PhD and two MS theses were completed, as well as two thesis proposals defended.

Acoustic modeling software improvements. In order to conduct distribution clustering experiments with more general features, such as speaking rate and lexical context, we implemented several major changes to our clustering system, fixing some bugs and updating libraries in the process. This effort also uncovered bugs in the adaptation work that may explain the fact that we were observing only small gains in performance with adaptation. Effort on this project is ongoing.

In addition, we explored improvements to our algorithm for estimating weights for score combination (i.e. LM and acoustic score weights used in rescoring). We developed a new and potentially more efficient way to estimate weights. Instead of evaluating points exhaustively on a very fine grid, as we do now, the aim is to avoid evaluating points which fall in the same "cell" (a polytope). Given a starting point in a cell, neighboring cells in the coordinate direction are evaluated for improvement to choose the new point, similar to steepest descent. Convergence is reached at a local minima. Experiments are underway to determine whether there are benefits to this approach.

Mixture language modeling. One of the important questions in language modeling today is how to effectively represent the long-term structure of language, i.e. how to capture dependence over longer sequences of words than can be modeled with a simple n-gram. To address this problem, we have developed a sentence-level mixture language model (LM) that represents the topic-dependent structure of language with separate n-gram language model mixture components determined using automatic clustering. In previous work, we obtained a 6% reduction in recognition error using a 5-component mixture models as compared to the standard trigram models on the 5k vocabulary WSJ H2 task. In this period we investigated two extensions – 1) development of a constrained

expectation-maximization algorithm (EM) for training the mixture models, and 2) introducing dynamic modeling into the sentence-level mixtures – as well as software changes to accommodate the new search algorithms.

The EM algorithm for training the component language models and their mixture weights is a relatively straightforward extension of general EM mixture model training. However, it is important to include some sort of back-off algorithm in estimating the n -gram probabilities. To include this in the EM training framework, we introduced a set of constraints analogous to the Witten-Bell back-off equations that set a minimum probability for the different n -grams. The EM algorithm gave a slight improvement in perplexity and recognition accuracy relative to the Viterbi-style training algorithm that we had implemented previously. In training on the 1994 multi-source LM training set, we observed a significant reduction in perplexity due to the mixture model, and noted that to some extent the component language models clustered by newspaper. However, this major gain in perplexity has not translated into improvements in recognition performance, a surprising result that we are trying to better understand.

Another approach to capturing topic-dependence is dynamic language modeling, which adjusts word frequencies depending on what words have been observed in the speech previously. Dynamic language model adaptation easily fits into the sentence-level mixture model framework in two ways. First, the sentence-level mixture weights can be recursively adapted according to the likelihood of the respective mixture components in the previous utterance, as in [12] for n -gram level mixture weights. Second, the dynamic n -gram cache model [10, 11] can be incorporated into the mixture language model. However, in the mixture model, it is possible to have component-dependent cache models, where each component cache would be updated after each sentence according to the likelihood of that component given the recognized word string. We also introduced new variations of cache language modeling, including a selective unigram cache including only content words and a word-class cache. This unigram cache was used alone and in addition to an interpolated bigram/trigram cache at the n -gram level. We conducted several supervised adaptation experiments based on the ARPA WSJ 1993 H2 (5k) development and evaluation data but were not able to show significant improvements because the article length was typically around three sentences long and in some cases not contiguous. We obtained an overall 3.5% reduction in word error rate and 11% reduction in perplexity with supervised dynamic adaptation.

During the past year, we have been developing a lattice-based decoder, as will be described in a subsequent section, which will serve as our baseline system in the future. The decoder uses a backward n -gram language model unlike our previous the N-best rescoring work, i.e. $p(w_i|w_{i+1}, w_{i+2})$ vs. $p(w_i|w_{i-1}, w_{i-2})$. Therefore, we implemented changes to our LM software to allow for training of and recognition with backwards language models.

Intra-utterance phoneme dependence modeling. Over the past year, we have developed the theoretical framework for a hierarchical model of dependence for a set of discrete random variables, which we plan to use as a model of intra-utterance phoneme dependence. We use a dependence tree [13] to represent the correlation among random variables, i.e. a tree structure (designed automatically) with Markov assumptions along the branches of the tree. The dependence tree can be thought of as representing a vector “state” that describes the speaker/utterance, where each element of the vector corresponds to a phoneme. Since most utterances will not contain all possible phonemes, we derived an efficient algorithm for computing the likelihood of the observed data, which we call the upward-downward algorithm to emphasize the analogy to the forward-backward algorithm. This algorithm, which we recently wrote up and submitted for publication [8], is needed for solving the parameter estimation problem when the tree structure is given, as the E-step in the EM algorithm. The EM algorithm is then used as one step in an iterative approach to combined dependence tree topology design and parameter estimation.

We have implemented the training algorithm for discrete distribution dependence trees (parameters and topology estimation), and have conducted initial experiments on the TIMIT corpus. These experiments raised the issue of distribution smoothing as an important problem that is not simply addressed by the effective smoothing of the EM algorithm. For the moment, we implemented some simple smoothing heuristics, but plan to explore this problem further. Additional future work is to extend the dependence tree design algorithm to include continuous distributions and to represent variable-length observations.

Channel estimation. In previous work, we evaluated two channel estimation algorithms in the context of the BU SSM recognition system on the WSJ S6 telephone task. In both cases, we estimated the channel compensation vector based on the full vector of cepstra and difference cepstra. Because a linear, time-invariant channel (which both models assume) would have a zero vector for compensating the difference cepstra, we tried estimation under this constraint, but no improvement in performance was observed.

Next, new channel compensation algorithms were developed, both based on using a prior channel model in the channel estimate. Two variations were developed: one that can be implemented in the feature space, i.e. subtracting the maximum a posteriori (MAP) channel estimate from the cepstral feature vectors, and one that requires modifications in the model space, i.e. shifting the mean and covariances of the triphone models to match a particular utterance using Bayesian learning. Initially, the channel prior information will be estimated from training data by finding a set of ML channel estimates and computing statistics from the estimates. Implementation of the algorithm is in progress. The Macrophone Natural Number corpus is being used as a test paradigm to narrow the focus of the task to short utterance, telephone speech recognition. HTK will be used to establish a baseline recognition result using a basic word-pair grammar, in order to reduce system building costs for this new task. [This work was supported by an ARPA AASERT award

associated with this project.]

Lattice search algorithms for multi-pass recognition scoring. In the last quarter, we implemented a lattice dynamic programming (DP) algorithm for rescoring HMM hypotheses, and demonstrated that it achieved comparable performance at a lower cost. Since then, we implemented a new local search algorithm that iteratively evaluates sentence level changes to the recognition hypothesis. (In addition, some improvements were made to the lattice representation, and a standard lattice file format for representing these lattices or any generic type of lattice was proposed and is being considered for use as a standard file format by the CSR community.) In recent work [14], we have been investigating the performance/speed trade-offs of three different fast lattice-based search algorithms: the lattice dynamic programming (DP) algorithm, a lattice N-best rescoring algorithm, and the lattice local search algorithm. In all cases, the goal is to find the sentence hypothesis with the highest combined score in a lattice of words. The lattice DP algorithm is an efficient optimal algorithm which guarantees that the highest scoring hypothesis will be found, but only Markov knowledge sources can be used with it. The lattice N-best and local search algorithms, on the other hand, allow incorporation of non-Markov models such as long-distance language models into the search. Of these two algorithms, the local search is sub-optimal but much faster.

Experiments were run on the WSJ 1993 5k word Hub 2 and 20k word Hub 1 tasks, using a combined score that included the BU SSM, the number of words, phones and inter-word silences and either a trigram LM or the BU sentence level mixture LM. We found that both the DP and local search algorithms attained comparable performance to N-best rescoring while running as much as 10 times faster. It was also demonstrated that the lattice local search algorithm had the advantage over the lattice DP algorithm of being able to use the BU sentence-level mixture LM, and therefore improve performance even though it is a sub-optimal search. We concluded that, for Markov knowledge sources, lattice DP is the most efficient search strategy and gives the best performance, but that the local search is better for incorporating sentence-level knowledge sources. The lattice N-best rescoring algorithm is still useful, however, for finding the scores of sentence hypotheses that are used in score combination weight estimation. Preliminary experiments on the Switchboard corpus showed similar or better gains in speed, but the overall error rates are too high to draw meaningful conclusions about performance. [This work was supported by an ONR AASERT award.]

ARPA benchmark tests. Significant effort during October and November went toward participation in the ARPA WSJ speech recognition benchmark tests. Unfortunately, because data for rescoring for development was available to us so late in the process, and because of a change in some file formats, we were unable to use many of our more recent system developments and therefore reported results on a system similar to that used last year. Under these circumstances, we were happy to achieve 11.6% word error, performance comparable to the results of many of the

other sites in the ARPA program. We plan to re-run a newer version of our system on this test in January.

The new work that we did include, the dynamic language model in the H1-C2 supervised adaptation contrast condition, did not give us the gains that we had hoped for, and we are looking more closely at the implementation and errors to understand this result.

Future Goals

The originally funded project comes to a close with this work, though research funded by the AASERT awards will continue in the areas of telephone channel estimation and search algorithms for long-distance dependence models. We are hoping for continued funding of this effort, where we have proposed to look at high-order modeling techniques for continuous speech recognition. In particular, we plan to concentrate on three problems: 1) hierarchical intra-utterance dependence modeling, extending the current work in this area; 2) unsupervised adaptation of acoustic models within and across utterances; and 3) sub-language modeling triggered by both acoustic and dialog-level cues. During the next six months, there will be a low-level effort on these problems, because Prof. Ostendorf will be on sabbatical at ATR in Japan (working in related areas) and many students will be on leave of absence working in industry.

References

- [1] M. Ostendorf and S. Roukos, "A Stochastic Segment Model for Phoneme-Based Continuous Speech Recognition," *IEEE Trans. Acoustics Speech and Signal Processing*, Dec. 1989.
- [2] S. Roukos, M. Ostendorf, H. Gish, and A. Derr, "Stochastic Segment Modeling Using the Estimate-Maximize Algorithm," *IEEE Int. Conf. Acoust., Speech, Signal Processing*, pages 127-130, New York, New York, April 1988.
- [3] M. Ostendorf, A. Kannan, S. Austin, O. Kimball, R. Schwartz, J. R. Rohlicek, "Integration of Diverse Recognition Methodologies Through Reevaluation of N-Best Sentence Hypotheses," *Proc. of the DARPA Workshop on Speech and Natural Language*, pp. 83-87, February 1991.
- [4] A. Kannan, M. Ostendorf and J. R. Rohlicek, "Maximum Likelihood Clustering of Gaussians for Speech Recognition," *IEEE Transactions on Speech and Audio Processing*, to appear.
- [5] A. Kannan and M. Ostendorf, "A Comparison of Trajectory and Mixture Modeling in Segment-based Word Recognition," *Proc. of the Inter. Conf. on Acoust., Speech and Signal Processing*, Vol. II, pp. 327-330, April 1993.

- [6] O. Kimball and M. Ostendorf, "On the Use of Tied Mixture Distributions," *Proceedings of the ARPA Workshop on Human Language Technology*, pp. 102-107, 1993.
- [7] O. Kimball, *Segment modeling alternatives for continuous speech recognition*, Ph.D. Thesis, E.C.S. Department, Boston University, 1994.
- [8] O. Ronen, J. R. Rohlicek and M. Ostendorf, "The Upward-Downward Algorithm for Computing Dependence Tree Likelihoods," manuscript submitted to *IEEE Signal Processing Letters*.
- [9] R. Iyer, M. Ostendorf and J. R. Rohlicek, "Language Modeling with Sentence-Level Mixtures," *Proceedings of the 1994 ARPA Workshop on Human Language Technology*, March 1994.
- [10] F. Jelinek, B. Merialdo, S. Roukos and M. Strauss, "A Dynamic LM for Speech Recognition", *Proc. DARPA Speech and Natural Language Workshop*, pp. 293-295, 1991.
- [11] R. Kuhn and R. de Mori, "A Cache Based Natural Language Model for Speech Recognition", *IEEE Transactions PAMI*, Vol. 14, pp. 570-583, 1992.
- [12] R. Kneser and V. Steinbiss, "On the Dynamic Adaptation of Stochastic LM", *Proc. of the Inter. Conf. on Acoust., Speech and Signal Processing*, Vol. 2, pp. 586-589, 1993.
- [13] C. K. Chow and C. N. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE Transactions on Information Theory*, Vol. IT-14, No. 3, May 1968, pp. 462-467.
- [14] F. Richardson, M. Ostendorf and J. R. Rohlicek, "Lattice-based Search Strategies for Large Vocabulary Recognition," to appear in *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 1995.

Principal Investigator Name: Mari Ostendorf

PI Institution: Boston University

PI Phone Number: 617-353-5430

PI E-mail Address: mo@raven.bu.edu

Grant or Contract Title: Segment-Based Acoustic Models for Continuous Speech Recognition

Grant or Contract Number: ONR-N00014-92-J-1778

Reporting Period: 1 July 1994 – 31 December 1994

3 Publications and Presentations

During this reporting period, we published one refereed paper, submitted an additional paper to a refereed journal, concluded three student theses, and Prof. Ostendorf gave one invited talk associated with this project, as itemized below. In addition, one Boston University M.S. thesis proposal was successfully defended by Becky Bates, entitled "Reducing the Effects of Telephone Channel Distortion and Additive Noise on Continuous Speech Recognition," and one Boston University Ph.D. thesis proposal was successfully defended by Orith Ronen, entitled "Hierarchical Models of Intra-Utterance Phoneme Dependence."

Refereed papers published:

"Maximum Likelihood Clustering of Gaussians for Speech Recognition," A. Kannan, M. Ostendorf and J. R. Rohlicek, *IEEE Trans. Speech and Audio Processing*, Vol. 2, No. 3, July 1994, pp. 453-455.

Refereed papers submitted but not yet published:

"The Upward-Downward Algorithm for Computing Dependence Tree Likelihoods," O. Ronen, J. R. Rohlicek and M. Ostendorf, manuscript submitted to *IEEE Signal Processing Letters*.

Unrefereed Reports and Conference Papers:

Segment Modeling Alternatives for Continuous Speech Recognition, O. Kimball, Boston University Ph.D. Thesis, 1994.

Language Modeling with Sentence-Level Mixtures, R. Iyer, Boston University M.S. Thesis, 1994.

Lattice-based Search Strategies for Large Vocabulary Speech Recognition, F. Richardson, Boston University M.S. Thesis, 1994.

Conference presentations and invited talks

"A Unified View of Stochastic Modeling for Speech Recognition", M. Ostendorf, invited talk at Johns Hopkins University, December 1994.

Principal Investigator Name: Mari Ostendorf
PI Institution: Boston University
PI Phone Number: 617-353-5430
PI E-mail Address: mo@raven.bu.edu
Grant or Contract Title: Segment-Based Acoustic Models for Continuous Speech Recognition
Grant or Contract Number: ONR-N00014-92-J-1778
Reporting Period: 1 July 1994 – 31 December 1994

4 Transitions and DoD Interactions

This grant includes a subcontract to BBN, and the research results and software is available to them. Thus far, we have collaborated with BBN by combining the Byblos system with the SSM in N-Best sentence rescoring to obtain improved recognition performance, and we have provided BBN with papers and technical reports to facilitate sharing of algorithmic improvements. On their part, BBN has been very helpful to us in our WSJ porting efforts, providing us with WSJ data and consulting on format changes.

We have also begun an effort to collaborate more closely in lattice rescoring. Boston University student Fred Richardson has implemented software libraries that will be shared by both sites, and he has modified the BBN decoder to provide lattices annotated with segmentation times and HMM scores.

The recognition system that has been developed under the support of this grant and of a joint NSF-ARPA grant (NSF # IRI-8902124) is currently being used for automatically obtaining good quality phonetic alignments for a corpus of radio news speech under development at Boston University in a project supported by the Linguistic Data Consortium.

Principal Investigator Name: Mari Ostendorf

PI Institution: Boston University

PI Phone Number: 617-353-5430

PI E-mail Address: mo@raven.bu.edu

Grant or Contract Title: Segment-Based Acoustic Models for Continuous Speech Recognition

Grant or Contract Number: ONR-N00014-92-J-1778

Reporting Period: 1 July 1994 – 31 December 1994

5 Software and Hardware Prototypes

Our research has required the development and refinement of software systems for parameter estimation and recognition search, which are implemented in C or C++ and run on Sun Sparc workstations. No commercialization is planned at this time.

January 14, 1995

Dear Family, HAPPY NEW YEAR!

Much has happened since my last letter, and I am anxious to report those events for all of you.

I suppose to bring you up to date with our travels we should start with our Tennessee trip to attend the wedding of Kathleen Sullivan, daughter of Sue and Bill Sullivan of San Jose Calif. We met the Turco's at the Nashville Airport and drove with them to the University of the South. If you have never heard of it, the Univ. of the South is in the town of Swanee-never heard of that either? The town is located about 80 miles Southeast of Nashville. The rehearsal dinner was held that night and we all attended that function. It is always interesting to meet the other families at such functions, because they are individuals you see then, at the wedding, the reception and probably never again. For some reason, that type of meeting always turns out well and everyone really enjoyed themselves, although I got a kick out of the local culture in their voices and they probably thought we also spoke with a strange dialect.

We stayed on the campus at a motel operated by the University, for the wedding was to take place at the University chapel. I say chapel, but it was anything but what your mind might conjure as a chapel. It is a huge edifice with breathtakingly beautiful stained glass windows. Polk Family history abounds in this part of the country. Bishop Leodontis Polk built this church. Kathleen and her father walked the entire length of the church with music coming from an organ with hundreds of pipes-it was a sound experience!. It was a small wedding party and guests, but it was quite an event. Following the ceremony everyone left for the Gore Family farm in a little town nearly 35 miles to the West. There the reception was held and a fabulous table of food was spread. The afternoon was truly beautiful with full sun, temperature in the 70's and not a cloud in the sky. Connie's sisters including Pat Dodd, Terry Turco, Sue Sullivan (mother of the bride) Pat Dodd's son John from Los Angeles, Mike Sullivan (brother of the bride) and the bride's sisters, Hilary, her husband and their three children and Sally, who lives in Burlingame were among those present.

The next day we visited the Polk mansion RATTLE AND SNAP, a National Monument. where Connie's grandfather Polk was born. The house is fantastic and has been restored both inside and outside. It majestically sits on a rise easily seen from the road and must be surrounded by at least 200 acres . From there we went into town to view President Polk's home. Following that it was lunch and off to the airport. Connie had called ahead and made arrangements for Marie (uncle Doc's daughter) and her husband Joe to meet us at the airport. We called them once we got in town and they met us for dinner. It was a lot of fun chatting with them. That was the first time Connie and I had met Joe- what a neat guy. I am sure all of you would enjoy his company. He is very serious about his music and obviously very industrious. Marie continues to pursue her M.S.studies and those should be completed before this time next year. The immediate hurdle, as I recall, is a practicum that must be completed.

For Christmas, we decided to go to Florida. We arranged the dates so as to arrive on Matthew's and Connie Jr's birthday, for as you will recall they were both born on the same date December 16-one year apart. Linda and Chris hosted the birthday party. What was so unique was the fact that our presence was a total surprise to the birthday two-that made it all the more enjoyable. Patty and her three children plus Ben and Andy, Baron, Meg, Chris, Billy and Matt Milana made for a large group.

We used our Condo as a launching pad to go to Pat's and Chris' in Ocala where we spent one night and we watched Christine and Carey take their horses through their paces as well as their jumps. Those two young girls are quite adept at riding, and it is clear they really enjoy the sport. To be convinced one has only to see their collection of trophy's. We made two trips to Sarasota, the first time to be with Connie Jr. and Baron and have a guided tour of their new home. It is great! Those two have a good start on life. There is plenty of room in that house for the little bundle of joy Connie is carrying that is due for appearance in August, 1995. The second trip to Sarasota was to have lunch with a couple that are spending the winter in Ft Myers and who live on our floor here at 3 Seal Harbor.

Most of the family attended Christmas Mass at St. Paul's and then came to our Condo the next morning about 11:00 a.m. for brunch and opening of

gifts from Santa. Lots of fun and laughter with Teresa and Joe, the Milanases, Ben and Matthew as well as Sue and her husband Allan, Tim, Jamie and Armanda, Connie Jr., Baron, Dawn and Andy.

Connie and I do not expect anything from our children for Christmas. We are satisfied that all our children and grandchildren are good, decent citizens who love each other and are genuinely happy to see us. At the Christmas morning gathering that there were small items that were presented to us- little did we know there was something in the wind!

That evening we had Christmas dinner at Linda and Chris' (29 in all)- turkey and ham and all the trimmings. WE all missed Liz and her children who were in Connecticut. Bill, Jocie and Callan arrived in time for that dinner having driven from Marietta Georgia that same day. After such a long drive, Callan seemed overwhelmed when she entered the house for there were so many people who professed to be uncles aunts etc. It was a great dinner, and as it began to come to a close, two rather large boxes appeared. One was placed in front of Connie and the other was placed in front of me. We had not the slightest hint of what to expect, and certainly did not in our wildest dream imagine what would be in those boxes. Upon removing the wrapping, what blew us out of the room were the contents- a Cam Corder!!!!-A great surprise that will certainly be used! Thanks to everyone!

Before Bill, Jocie and Callan left town, Matthew invited us to dinner and served Brouts, boiled in beer and browned on the grill. If you have not tried Brouts, you owe it to yourself. Fat grams be darned!.

While we were in sunny Tampa, there were winds in excess of 80 mph and driving rains at the Winthrop address. When we returned there was a little water damage inside our condo-certainly is nice to be a renter! As you might guess, the first thing we did was turn on the heaters. The temperature in the Condo was 51 degrees and it took about a day to raise that temp to something that was comfortable. It was a tough battle because the outside temp with the wind-chill was -18 and the actual temp was +13 degrees. It was then I decided that we needed some additional insulation. We now have covered all the windows with plastic and plugged all the holes that was permitting the cold air to enter. The next time old man winter throws himself at us we expect to be prepared. I told Connie that plastic is there for the duration. In this case the duration is the first full moon in May-that is when

the farmers plant crops that are vulnerable to freezing temperatures in New England. Connie is most unhappy being wrapped in plastic. In fact, she is now calling the condo a bunker! Sound anxious to get out?

I am pleased to report that all seems well in the extended Taft family. We are both looking forward to August, 1996, for that is when we will pack it up and head for Florida, never more to put up with what these natives call the four seasons. From our perspective, after living in Florida from 1963 until 1985, there are two seasons in new England. they are June July and August when one can go to work in short sleeve shirts, a light sport coat and no wool pants. The others season is long sleeve shirts, wool sport coat wool stockings and wool pants. That reminds me, when Connie and I emerged from the plane in Tampa on December, you should have seen us peeling off the wool sweaters and jackets. It had been 18 degrees when we left Boston and must have been in the 80's in Tampa. How the weather can change in three hours (of flying time)

Connie and I wish each of you a happy, healthy prosperous and holy 1995.

Love to each and everyone

Mom and Dad.